arXiv:1702.02397v1 [cs.LO] 8 Feb 2017

Quasi-reductivity of Logically Constrained Term Rewriting Systems

CYNTHIA KOP, University of Innsbruck and University of Copenhagen

This paper considers *quasi-reductivity*—essentially, the property that an evaluation cannot get "stuck" due to a missing case in pattern matching—in the context of term rewriting with logical constraints.

1. INTRODUCTION

The formal framework of *Logically Constrained Term Rewriting Systems* (LCTRSs), introduced in [Kop and Nishida 2013], combines term rewriting with constraints and calculations over an arbitrary theory. This for instance allows users to specify rules with integers, arrays and strings, and can be used to analyze both imperative and functional programs (without higher-order variables) in a natural way.

Many methods to analyze term rewriting systems naturally extend to LCTRSs. In this paper we will study *quasi-reductivity*, the property that the only ground irreducible terms are constructor terms. We provide a simple method to prove quasi-reductivity: essentially, we will test that the rules do not omit any patterns.

Structure: For completeness, we will first set out the definition of LCTRSs, following [Kop and Nishida 2013; Kop and Nishida 2014] in Section 2. In Section 3 we consider the definition of quasi-reductivity; in Section 4 we present three restrictions: *left-linearity*, *constructor-soundness* and *left-value-freeness*. The core of this work is Section 5, where we provide an algorithm to confirm quasi-reductivity for LCTRSs which satisfy these restrictions, and prove its soundness.

2. PRELIMINARIES

In this section, we briefly recall *Logically Constrained Term Rewriting Systems* (usually abbreviated as *LCTRSs*), following the definitions in [Kop and Nishida 2013].

Many-sorted terms. We introduce terms, typing, substitutions, contexts and subterms (with corresponding terminology) in the usual way for many-sorted term rewriting.

Definition 2.1. We assume given a set $\mathcal{S}$ of *sorts* and an infinite set $\mathcal{V}$ of *variables*, each variable equipped with a sort. A *signature* Σ is a set of *function symbols* f, disjoint from $\mathcal{V}$, each equipped with a *sort declaration* $[\iota_1 \times \cdots \times \iota_n] \Rightarrow \kappa$, with all ι_i and κ sorts. For readability, we often write κ instead of $[] \Rightarrow \kappa$. The set $\mathcal{T}erms(\Sigma, \mathcal{V})$ of *terms* over Σ and $\mathcal{V}$ contains any expression s such that $\vdash s : \iota$ can be derived for some sort ι, using:

$$\frac{}{\vdash x : \iota}\ (x : \iota \in \mathcal{V}) \qquad \frac{\vdash s_1 : \iota_1 \quad \ldots \quad \vdash s_n : \iota_n}{\vdash f(s_1, \ldots, s_n) : \kappa}\ (f : [\iota_1 \times \cdots \times \iota_n] \Rightarrow \kappa \in \Sigma)$$

We fix Σ and $\mathcal{V}$. Note that for every term s, there is a unique sort ι with $\vdash s : \iota$.

Definition 2.2. Let $\vdash s : \iota$. We call ι the *sort of* s. Let $Var(s)$ be the set of variables occurring in s; we say that s is *ground* if $Var(s) = \emptyset$.

Definition 2.3. A *substitution* γ is a sort-preserving total mapping from $\mathcal{V}$ to $\mathcal{T}erms(\Sigma, \mathcal{V})$. The result $s\gamma$ of applying a substitution γ to a term s is s with all occurrences of a variable x replaced by $\gamma(x)$. The *domain* of this substitution, $Dom(\gamma)$, is the set of variables x with $\gamma(x) \neq x$. The notation $[x_1 := s_1, \ldots, x_k := s_k]$ denotes a substitution γ with $\gamma(x_i) = s_i$ for $1 \leq i \leq n$, and $\gamma(y) = y$ for $y \notin \{x_1, \ldots, x_n\}$.

Definition 2.4. A *context* C is a term containing a typed *hole* $\Box : \iota$. If $t : \iota$, we define $C[t_1]$ as C with $\Box$ replaced by t. If we can write $s = C[t]$, then t is a *subterm* of s.

This work is supported by the Austrian Science Fund (FWF) international project I963, and the Marie-Skłodowska-Curie action "HORIP", program H2020-MSCA-IF-2014, 658162.

Logical terms. Specific to LCTRSs, we consider different kinds of symbols and terms.

Definition 2.5. We assume given:

— signatures Σ_{terms} and Σ_{theory} such that $\Sigma = \Sigma_{terms} \cup \Sigma_{theory}$;
— a mapping $\mathcal{I}$ which assigns to each sort ι occurring in Σ_{theory} a set $\mathcal{I}_\iota$;
— a mapping $\mathcal{J}$ which assigns to each $f : [\iota_1 \times \cdots \times \iota_n] \Rightarrow \kappa \in \Sigma_{theory}$ a function in $\mathcal{I}_{\iota_1} \times \cdots \times \mathcal{I}_{\iota_n} \Longrightarrow \mathcal{I}_\kappa$;
— for all sorts ι occurring in Σ_{theory} a set $\mathcal{V}al_\iota \subseteq \Sigma_{theory}$ of *values*: function symbols $a : [] \Rightarrow \iota$ such that $\mathcal{J}$ gives a bijective mapping from $\mathcal{V}al_\iota$ to $\mathcal{I}_\iota$.

We require that $\Sigma_{terms} \cap \Sigma_{theory} \subseteq \mathcal{V}al = \bigcup_\iota \mathcal{V}al_\iota$. The sorts occurring in Σ_{theory} are called *theory sorts*, and the symbols *theory symbols*. Symbols in $\Sigma_{theory} \setminus \mathcal{V}al$ are *calculation symbols*. A term in $\mathcal{T}erms(\Sigma_{theory}, \mathcal{V})$ is called a *logical term*.

Definition 2.6. For ground logical terms, let $[\![f(s_1, \ldots, s_n)]\!] := \mathcal{J}_f([\![s_1]\!], \ldots, [\![s_n]\!])$. Every ground logical term s corresponds to a unique value c such that $[\![s]\!] = [\![c]\!]$; we say that c is the value of s. A *constraint* is a logical term φ of some sort bool with $\mathcal{I}_{\mathsf{bool}} = \mathbb{B} = \{\top, \bot\}$, the set of *booleans*. A constraint φ is *valid* if $[\![\varphi\gamma]\!] = \top$ for *all* substitutions γ which map $Var(\varphi)$ to values, and *satisfiable* if $[\![\varphi\gamma]\!] = \top$ for *some* substitutions γ which map $Var(\varphi)$ to values. A substitution γ *respects* φ if $\gamma(x)$ is a value for all $x \in Var(\varphi)$ and $[\![\varphi\gamma]\!] = \top$.

Formally, terms in $\mathcal{T}erms(\Sigma_{terms}, \mathcal{V})$ have no special function, but we see them as the primary objects of our term rewriting systems: a reduction would typically begin and end with such terms, with calculation symbols only used in intermediate terms. Their function is to perform calculations in the underlying theory. Usually, values which are expected to occur in starting terms and end terms should be included both in Σ_{terms} and Σ_{theory}, while values only used in constraints and calculations would only be in Σ_{theory}; true and false often fall in the latter category.

Example 2.7. Let $\mathcal{S} = \{\mathsf{int}, \mathsf{bool}\}$, and consider the signature $\Sigma = \Sigma_{terms} \cup \Sigma_{theory}$ where $\Sigma_{terms} = \{\ \mathsf{fact} : [\mathsf{int}] \Rightarrow \mathsf{int}\ \} \cup \{\ \mathsf{n} : \mathsf{int} \mid n \in \mathbb{Z}\ \}$ and $\Sigma_{theory} = \{\ \mathsf{true}, \mathsf{false} : \mathsf{bool}, \wedge, \vee, \Rightarrow : [\mathsf{bool} \times \mathsf{bool}] \Rightarrow \mathsf{bool},\ +, -, * : [\mathsf{int} \times \mathsf{int}] \Rightarrow \mathsf{int},\ \leq, <, = : [\mathsf{int} \times \mathsf{int}] \Rightarrow \mathsf{bool}\ \} \cup \{\mathsf{n} : \mathsf{int} \mid n \in \mathbb{Z}\}$. Then both int and bool are theory sorts, and the values are true, false and all symbols n representing integers. For the interpretations, let $\mathcal{I}_{\mathsf{int}} = \mathbb{Z}$, $\mathcal{I}_{\mathsf{bool}} = \mathbb{B}$, and let $\mathcal{J}$ be the evaluation function which interprets these symbols as expected.

Using infix notation, examples of logical terms are $0 = 0 + -1$ and $x + 3 \leq y + -42$. Both are constraints. $5{+}9$ is also a (ground) logical term, but not a constraint. Expected starting terms are for instance $\mathsf{fact}(42)$ or $\mathsf{fact}(\mathsf{fact}(-4))$: ground terms fully built using symbols in Σ_{terms}.

Rules and rewriting. We adapt the standard notions of rewriting (see, e.g., [Baader and Nipkow 1998]) by including constraints and adding rules to perform calculations.

Definition 2.8. A *rule* is a triple $\ell \to r\ [\varphi]$ where ℓ and r are terms of the same sort and φ is a constraint. Here, ℓ has the form $f(\ell_1, \ldots, \ell_n)$ and contains at least one symbol in $\Sigma_{terms} \setminus \Sigma_{theory}$ (so ℓ is not a logical term). If $\varphi = \mathsf{true}$ with $\mathcal{J}(\mathsf{true}) = \top$, the rule may be denoted $\ell \to r$. Let $LVar(\ell \to r\ [\varphi])$ denote $Var(\varphi) \cup (\ Var(r) \setminus Var(\ell))$. A substitution γ *respects* $\ell \to r\ [\varphi]$ if $\gamma(x)$ is a value for all $x \in LVar(\ell \to r\ [\varphi])$, and $[\![\varphi\gamma]\!] = \top$.

Note that it is allowed to have $Var(r) \not\subseteq Var(\ell)$, but fresh variables in the right-hand side may only be instantiated with *values*. This is done to model user input or random choice, both of which would typically produce a value. Variables in the left-hand sides do not need to be instantiated with values (unless they also occur in the constraint); this is needed for instance to support a lazy evaluation strategy.

Definition 2.9. We assume given a set of rules $\mathcal{R}$ and let $\mathcal{R}_{\mathtt{calc}}$ be the set $\{f(x_1,\ldots,x_n) \to y\ [y = f(\overrightarrow{x})] \mid f : [\iota_1 \times \cdots \times \iota_n] \Rightarrow \kappa \in \Sigma_{theory} \setminus \mathcal{V}al\}$ (writing $\overrightarrow{x}$ for $x_1,\ldots,x_n$). The *rewrite relation* $\to_{\mathcal{R}}$ is a binary relation on terms, defined by:

$$C[\ell\gamma] \to_{\mathcal{R}} C[r\gamma] \quad \text{if } \ell \to r\ [\varphi] \in \mathcal{R} \cup \mathcal{R}_{\mathtt{calc}} \text{ and } \gamma \text{ respects } \ell \to r\ [\varphi]$$

Here, C is an arbitrary context. A reduction step with $\mathcal{R}_{\mathtt{calc}}$ is called a *calculation*. A term is in *normal form* if it cannot be reduced with $\to_{\mathcal{R}}$.

We will usually call the elements of $\mathcal{R}_{\mathtt{calc}}$ rules—or *calculation rules*–even though their left-hand side is a logical term.

Definition 2.10. For $f(\ell_1,\ldots,\ell_n) \to r\ [\varphi] \in \mathcal{R}$ we call f a *defined symbol*; non-defined elements of Σ_{terms} and all values are *constructors*. Let $\mathcal{D}$ be the set of all defined symbols, and $\mathcal{C}ons$ the set of constructors. A term in $\mathcal{T}erms(\mathcal{C}ons, \mathcal{V})$ is a *constructor term*.

Now we may define a *logically constrained term rewriting system* (LCTRS) as the abstract rewriting system $(\mathcal{T}erms(\Sigma, \mathcal{V}), \to_{\mathcal{R}})$. An LCTRS is usually given by supplying Σ, $\mathcal{R}$, and also $\mathcal{I}$ and $\mathcal{J}$ if these are not clear from context.

Example 2.11. To implement an LCTRS calculating the *factorial* function, we use the signature Σ from Example 2.7, and the following rules:

$$\mathcal{R}_{\mathsf{fact}} = \{\ \mathsf{fact}(x) \to 1\ [x \leq 0]\ \ ,\ \ \mathsf{fact}(x) \to x * \mathsf{fact}(x-1)\ [\neg(x \leq 0)]\ \}$$

Using calculation steps, a term $3-1$ reduces to 2 in one step (using the calculation rule $x - y \to z\ [z = x - y]$), and $3 * (2 * (1 * 1))$ reduces to 6 in three steps. Using also the rules in $\mathcal{R}_{\mathsf{fact}}$, $\mathsf{fact}(3)$ reduces in ten steps to 6.

Example 2.12. To implement an LCTRS calculating the sum of elements in an array, let $\mathcal{I}_{\mathsf{bool}} = \mathbb{B}$, $\mathcal{I}_{\mathsf{int}} = \mathbb{Z}$, $\mathcal{I}_{\mathsf{array(int)}} = \mathbb{Z}^*$, so $\mathsf{array(int)}$ is mapped to finite-length integer sequences. Let $\Sigma_{theory} = \Sigma^{int}_{theory} \cup \{\mathsf{size} : [\mathsf{array(int)}] \Rightarrow \mathsf{int},\ \mathsf{select} : [\mathsf{array(int)} \times \mathsf{int}] \Rightarrow \mathsf{int}\} \cup \{\mathsf{a} \mid a \in \mathbb{Z}^*\}$. (We do *not* encode arrays as lists: every "array"—integer sequence—a corresponds to a unique symbol a.) The interpretation function $\mathcal{J}$ behaves on Σ^{int}_{theory} as usual, maps the values a to the corresponding integer sequence, and has:

$$\begin{array}{rcll} \mathcal{J}_{\mathsf{size}}(a) & = & k & \text{if } a = \langle n_0,\ldots,n_{k-1}\rangle \\ \mathcal{J}_{\mathsf{select}}(a,i) & = & n_i & \text{if } a = \langle n_0,\ldots,n_{k-1}\rangle \text{ with } 0 \leq i < k \\ & & 0 & \text{otherwise} \end{array}$$

In addition, let:

$$\begin{array}{rcl} \Sigma_{terms} & = & \{\ \mathsf{sum} : [\mathsf{array(int)}] \Rightarrow \mathsf{int},\ \mathsf{sum0} : [\mathsf{array(int)} \times \mathsf{int}] \Rightarrow \mathsf{int}\ \} \cup \\ & & \{\ \mathsf{n} : \mathsf{int} \mid n \in \mathbb{Z}\ \} \cup \{\ \mathsf{a} \mid a \in \mathbb{Z}^*\ \} \end{array}$$

$$\mathcal{R} = \left\{ \begin{array}{rcll} \mathsf{sum}(x) & \to & \mathsf{sum0}(x, \mathsf{size}(x) - 1) & \\ \mathsf{sum0}(x,k) & \to & \mathsf{select}(x,k) + \mathsf{sum0}(x,k-1) & [k \geq 0] \\ \mathsf{sum0}(x,k) & \to & 0 & [k < 0] \end{array} \right\}$$

Note the special role of *values*, which are new in LCTRSs compared to older styles of constrained rewriting. Values are the representatives of the underlying theory. All values are constants (constructor symbols $v()$ which do not take arguments), even if they represent complex structures, as seen in Example 2.12. However, not all constants are values; for instance a constant constructor $\mathsf{error} \in \Sigma_{terms}$ would not be a value. We will often work with signatures having infinitely many values. Note that we do not match modulo theories, e.g. we do not equate $0 + (x + y)$ with $y + x$ for matching.

Note also the restriction on variables in a constraint being instantiated by values; for instance in Example 2.11, a term $\mathsf{fact}(\mathsf{fact}(3))$ reduces only at the *inner* fact.

3. QUASI-REDUCTIVITY

The most high-level definition of quasi-reductivity is likely the following.

Definition 3.1 (*Quasi-reductivity*). An LCTRS $(\Sigma_{terms}, \Sigma_{theory}, \mathcal{I}, \mathcal{J}, \mathcal{R})$ is quasi-reductive if for all $s \in \mathcal{T}erms(\Sigma, \emptyset)$ one of the following holds:

— $s \in \mathcal{T}erms(\mathcal{C}ons, \emptyset)$ (we say: s is a *ground constructor term*);
— there is a t such that $s \to_{\mathcal{R}} t$ (we say: s *reduces*).

Note that $\mathcal{T}erms(\Sigma, \emptyset)$ is the set of *ground* terms. Another common definition concerns only the reduction of "basic" ground terms, but is equivalent:

LEMMA 3.2. *An LCTRS is quasi-reductive if and only if all terms $f(s_1, \ldots, s_n)$ with f a defined or calculation symbol and all $s_i \in \mathcal{T}erms(\mathcal{C}ons, \emptyset)$, reduce.*

PROOF. If the LCTRS is quasi-reductive, then each such $f(\overrightarrow{s})$ reduces, as it is not a constructor term. If the LCTRS is not quasi-reductive, then let $f(\overrightarrow{s})$ be a minimal ground irreducible non-constructor term. By minimality, all s_i must be constructor terms. If f is a constructor, then the whole term is a constructor term, contradiction, so f is either a defined symbol or a calculation symbol. □

4. RESTRICTIONS

For our algorithm in the next section, which proves that a given LCTRS is quasi-reductive, we will limit interest to LCTRSs which satisfy the following restrictions:

Definition 4.1 (*Restrictions*). An LCTRS $(\Sigma_{terms}, \Sigma_{theory}, \mathcal{I}, \mathcal{J}, \mathcal{R})$ is:

— *left-linear* if for all rules $\ell \to r\ [\varphi] \in \mathcal{R}$: every variable in ℓ occurs only once;
— *constructor-sound* if there are ground constructor terms for every sort ι such that some $f : [\ldots \times \iota \times \ldots] \Rightarrow \kappa \in \mathcal{D}$ (so for every input sort of a defined symbol);
— *left-value-free* if the left-hand sides of rules do not contain any values.

Note that any LCTRS can be turned into a left-value-free one, by replacing a value v by a fresh variable and adding a constraint $x = v$ instead. Constructor-soundness seems quite natural, with a sort representing the set of ground constructor terms of that sort. Left-linearity is probably the greatest limitation; however, note that non-left-linear systems impose *syntactic* equality. In a rule

$$\mathsf{addtoset}(x, \mathsf{setof}(x, rest)) \to \mathsf{setof}(x, rest)$$

we can reduce $\mathsf{addtoset}(3+4, \mathsf{setof}(3+4, s))$ immediately to $\mathsf{setof}(3+4, s)$. However, we cannot reduce $\mathsf{addtoset}(3+4, \mathsf{setof}(4+3, s))$ with this rule. There is also no syntactic way to check for *inequality*. Therefore, it seems like we could better formulate this rule and its complement using constraints, or (if the sort of x has non-value constructors) by a structural check. The rule above and its complement could for instance become:

$$\begin{array}{lll} \mathsf{addtoset}(x, \mathsf{setof}(y, rest)) & \to\ \mathsf{setof}(y, rest) & [x = y] \\ \mathsf{addtoset}(x, \mathsf{setof}(y, rest)) & \to\ \mathsf{setof}(y, \mathsf{addtoset}(x, rest)) & [x \neq y] \end{array}$$

In this light, left-linearity also seems like a very natural restriction.

Comment: In [Falke and Kapur 2012] a similar method is introduced to prove quasi-reductivity of a different style of constrained rewriting. There, however, the systems are additionally restricted to be *value-safe*: the only constructors of sorts occurring in Σ_{theory} are values. We drop this requirement here, because it is not necessary in the definition of LCTRSs.

Constructor-soundness, arguably the most innocent of these restrictions, allows us to limit interest to certain well-behaved rules when proving quasi-reductivity:

THEOREM 4.2. *A constructor-sound LCTRS with rules $\mathcal{R}$ is quasi-reductive if and only if the following conditions both hold:*

— *the same LCTRS restricted to constructor rules $\mathcal{R}' := \{f(\overrightarrow{\ell}) \to r\ [\varphi] \in \mathcal{R} \mid \forall i(\ell_i \in \mathcal{T}erms(\mathcal{C}ons, \mathcal{V}))\}$ is quasi-reductive;*
— *all constructor symbols with respect to $\mathcal{R}'$ are also constructors w.r.t. $\mathcal{R}$.*

PROOF. Suppose $\mathcal{R}'$ is quasi-reductive, and constructor terms are the same in either LCTRS. Then also $\mathcal{R}$ is quasi-reductive, as anything which reduces under $\mathcal{R}'$ also reduces under $\mathcal{R}$. Alternatively, suppose $\mathcal{R}$ is quasi-reductive.

Towards a contradiction, suppose $\mathcal{R}'$ has constructor symbols which are not constructors in $\mathcal{R}$; let f be such a symbol. As f is a constructor for $\mathcal{R}'$, there are no rules $f(\overrightarrow{l}) \to r\ [\varphi] \in \mathcal{R} \cup \mathcal{R}_{\mathtt{calc}}$ which match terms of the form $f(\overrightarrow{s})$ with all $s_i \in \mathcal{T}erms(\mathcal{C}ons, \emptyset)$. Because f is a defined symbol, such terms exist by constructor-soundness. As nothing matches $f(\overrightarrow{s})$ itself, and its strict subterms are constructor terms so cannot be reduced, this term contradicts quasi-reductivity of $\mathcal{R}$!

For the first point, suppose towards a contradiction that $\mathcal{R}'$ is not quasi-reductive, yet $\mathcal{R}$ is, and the same terms are constructor terms in either. By Lemma 3.2 there is some irreducible $f(s_1, \ldots, s_n)$ with all s_i constructor terms and f not a constructor. As the s_i are constructor terms, the rules in $\mathcal{R} \setminus \mathcal{R}'$ also cannot match! Thus, the term is also irreducible with $\to_{\mathcal{R}}$, contradiction. □

5. AN ALGORITHM TO PROVE QUASI-REDUCTIVITY

We now present an algorithm to confirm quasi-reductivity of a given LCTRS satisfying the restrictions from Definition 4.1. Following Theorem 4.2, we can—without loss of generality—limit interest to *constructor TRSs*, where the immediate arguments in the left-hand sides of rules are all constructor terms.

Main Algorithm. We assume given sequences $\iota_1, \ldots, \iota_n$ of theory sorts, $\kappa_1, \ldots, \kappa_m$ of sorts, and $x_1, \ldots, x_n$ of variables, with each $x_i : \iota_i \in \mathcal{V}$. Moreover, we assume given a set A of pairs $(\overrightarrow{s}, \varphi)$. Here, $\overrightarrow{s}$ is a sequence $s_1, \ldots, s_m$ of constructor terms which do not contain values, such that $\vdash s_i : \kappa_i$, and φ is a logical constraint. The s_i have no overlapping variables with each other or the x_j; that is, a term $f(x_1, \ldots, x_n, s_1, \ldots, s_m)$ would be linear. Variables in $\overrightarrow{x}$ and $\overrightarrow{s}$ may occur in φ, however.

Now, for $b \in \{\mathsf{term}, \mathsf{value}, \mathsf{either}\}$,[1] define the function $\boldsymbol{OK}(\overrightarrow{x}, A, b)$ as follows; this construction is well-defined by induction first on the number of function symbols occurring in A, second by the number of variables occurring in A, and third by the flag b (with $\mathsf{either} > \mathsf{term}, \mathsf{value}$). Only symbols in the terms s_i are counted for the first induction hypothesis, so not those in the constraints.

— if $m = 0$: let $\{y_1, \ldots, y_k\} = (\bigcup_{((),\varphi)\in A} Var(\varphi)) \setminus \{x_1, \ldots, x_n\}$;
  — if $\exists y_1 \ldots y_k(\bigvee_{((),\varphi)\in A} \varphi)$ is valid, then **true**
  — else **false**
  Note that if $A = \emptyset$, this returns **false**.
— if $m > 0$ and $b = \mathsf{either}$, then consider κ_1. If κ_1 does not occur in Σ_{theory}, the result is:

$$\boldsymbol{OK}(\overrightarrow{x}, A, \mathsf{term})$$

If κ_1 occurs in Σ_{theory} and all constructors with output sort κ_1 are values, then the result is:

$$\boldsymbol{OK}(\overrightarrow{x}, A, \mathsf{value})$$

[1] This parameter indicates what constructor instantiations we should consider for s_1.

If κ_1 occurs in Σ_{theory} but there are also non-value constructors of sort κ_1, then let $V := \{(\overrightarrow{s}, \varphi) \in A \mid s_1 \text{ is a variable}\}$ and $T := \{(\overrightarrow{s}, \varphi) \in A \mid s_1 \text{ is not a variable in } Var(\varphi)\}$. Note that V and T overlap in cases where s_1 is a variable not occurring in φ. The result of the function is:

$$\boldsymbol{OK}(\overrightarrow{x}, V, \mathsf{value}) \wedge \boldsymbol{OK}(\overrightarrow{x}, T, \mathsf{term})$$

In all cases, the recursive calls are defined, by the decrease in the third argument (and in the last case possibly also in the first and second argument).

— if $m > 0$ and $b = \mathsf{value}$, then we assume that κ_1 occurs in Σ_{theory} and for all $(\overrightarrow{s}, \varphi) \in A$ the first term, s_1, is a variable (if not, we might define the function result as **false**, but this cannot occur in the algorithm). Then let x_{n+1} be a fresh variable of sort κ_1, and let $A' := \{((s_2, \ldots, s_m), \varphi[s_1 := x_{n+1}]) \mid ((s_1, \ldots, s_n), \varphi) \in A\}$; the result is:

$$\boldsymbol{OK}((x_1, \ldots, x_{n+1}), A', \mathsf{either})$$

Note that A' has equally many function symbols as and fewer variables than A, and that we indeed have suitable sort sequences $(\iota_1, \ldots, \iota_n, \kappa_1)$ for the variables and $\kappa_2, \ldots, \kappa_m$ for the term sequences);

— if $m > 0$ and $b = \mathsf{term}$ and for all $(\overrightarrow{s}, \varphi) \in A$ the first term, s_1, is a variable, then we assume (like we did in the previous case) that never $s_1 \in \varphi$, and let $A' := \{((s_2, \ldots, s_m), \varphi) \mid ((s_1, \ldots, s_m), \varphi) \in A\}$; the result is:

$$\boldsymbol{OK}(\overrightarrow{x}, A', \mathsf{either})$$

A' has at most as many function symbols as and fewer variables than A.

— if $m > 0$ and $b = \mathsf{term}$ and there is some $(\overrightarrow{s}, \varphi) \in A$ where s_1 is not a variable, then let $f_1, \ldots, f_k$ be all non-value constructors with output sort κ_1 and let $A_1, \ldots, A_k$ be defined as follows: $A_i := \{(\overrightarrow{s}, \varphi) \in A \mid s_1 \text{ is a variable or has the form } f_i(\overrightarrow{t})\}$.
Now, for all i: if f_i has sort declaration $[\mu_1 \times \cdots \times \mu_p] \Rightarrow \kappa_1$, then we consider the new sort sequence $\overrightarrow{\kappa'}$ with $\overrightarrow{\kappa'} = (\mu_1, \ldots, \mu_p, \kappa_2, \ldots, \kappa_m)$; for every $(\overrightarrow{s}, \varphi) \in A_i$ we define:
  — if $s_1 = f_i(t_1, \ldots, t_p)$, then $\overrightarrow{u} := (t_1, \ldots, t_p, s_2, \ldots, s_m)$ and $\psi := \varphi$
  — if s_1 is a variable, then let $y_1, \ldots, y_p$ be fresh variables with sorts $\mu_1, \ldots, \mu_p$ respectively, and let $\overrightarrow{u} := (y_1, \ldots, y_p, s_2, \ldots, s_m)$ and $\psi := \varphi[s_1 := f_i(y_1, \ldots, y_p)]$.
We let B_i be the set of the corresponding $(\overrightarrow{u}, \psi)$ for all $(\overrightarrow{s}, \varphi) \in A_i$. Now, if A_i contains any element where s_i is not a variable, then B_i contains fewer function symbols (not counting constraints), so also fewer than A (as $A_i \subseteq A$). If all s_i are variables, then A_i is a strict subset of A, which misses at least one element where s_i contains a symbol, so also B_i has fewer symbols. Either way, we are safe defining the result as:

$$\boldsymbol{OK}(\overrightarrow{x}, B_1, \mathsf{either}) \wedge \cdots \wedge \boldsymbol{OK}(\overrightarrow{x}, B_k, \mathsf{either})$$

Correctness of this algorithm is proved using the following technical result.

LEMMA 5.1. *For any suitable $n, m, \overrightarrow{\iota}, \overrightarrow{\kappa}, \overrightarrow{x}, b$ and A such that $\boldsymbol{OK}(\overrightarrow{x}, A, b) = \mathtt{true}$, we have, for any sequence $(s_1, \ldots, s_n)$ of values and any sequence $(t_1, \ldots, t_m)$ of ground constructor terms: if one of the following conditions holds,*

— $b = \mathsf{either}$ *or* $m = 0$
— $b = \mathsf{value}$ *and* t_1 *is a value*
— $b = \mathsf{term}$ *and* t_1 *is not a value, and for all* $((u_1, \ldots, u_n), \varphi) \in A$*:* $u_1 \notin Var(\varphi)$

then there is some $((u_1, \ldots, u_m), \varphi) \in A$ *and a substitution* γ *with* $\gamma(x_i) = s_i$ *such that:*

— *each* $t_i = u_i\gamma$
— $\varphi\gamma$ *is a valid ground logical constraint*

PROOF. By induction on the derivation of $\boldsymbol{OK}(\overrightarrow{x}, A, b) = \mathtt{true}$. Let values $(s_1, \ldots, s_n)$ and ground constructor terms $(t_1, \ldots, t_m)$ which satisfy the conditions be given.

If $m = 0$, then let $\psi := \bigvee_{((),\varphi)\in A} \varphi$. By definition of $\boldsymbol{OK}$, $\exists y_1 \ldots y_k(\psi)$ is valid and $Var(\psi) \subseteq \{x_1, \ldots, x_n, y_1, \ldots, y_k\}$. That is, there are values $v_1, \ldots, v_k$ such that for all values $u_1, \ldots, u_n$ the ground constraint $\psi[\overrightarrow{y} := \overrightarrow{v}, \overrightarrow{x} := \overrightarrow{u}]$ is valid. In particular, we can take $\overrightarrow{s}$ for $\overrightarrow{u}$. Define $\gamma := [y_1 := v_1, \ldots, y_k := v_k, x_1 := s_1, \ldots, x_n := s_n]$. Then $\psi\gamma$ is valid, and since it is ground, some clause in the disjunction must be valid; so some $((), \varphi) \in A$ where $\varphi\gamma$ is a valid ground constraint. This is what the lemma requires.

If $m > 0$ and $b = \mathsf{either}$ and κ_1 does not occur in Σ_{theory}, then $\boldsymbol{OK}(\overrightarrow{x}, A, \mathsf{term})$ holds. Since there are no values of sort κ_1, the term t_1 is not a value; for the same reason, variables of sort κ_1 cannot occur in any constraint φ. Thus, the conditions for the induction hypothesis are satisfied; we find a suitable γ and $(\overrightarrow{u}, \varphi)$.

If $m > 0$ and $b = \mathsf{either}$ and κ_1 does occur in Σ_{theory}, and all constructors with output sort κ_1 are values, then $\boldsymbol{OK}(\overrightarrow{x}, A, \mathsf{value})$ holds; moreover, t_1 is necessarily a value, so we can again apply the induction hypothesis.

If $m > 0$ and $b = \mathsf{either}$ and there are both values and other constructors with output sort κ_1, then both $\boldsymbol{OK}(\overrightarrow{x}, T, \mathsf{term})$ and $\boldsymbol{OK}(\overrightarrow{x}, V, \mathsf{value})$ must hold. If t_1 is a value, then the conditions to apply the induction hypothesis with V are satisfied; if not, the conditions to apply it with T are satisfied! Since both T and V are subsets of A, this results in a suitable element and substitution.

If $m > 0$ and $b = \mathsf{value}$, then we may assume that t_1 is a value. Let x_{n+1} be a fresh variable of sort κ_1, and $A' := \{((u_2, \ldots, u_m), \varphi[u_1 := x_{n+1}]) \mid (\overrightarrow{u}, \varphi) \in A\}$. Applying the induction hypothesis with $n+1, m-1, (\overrightarrow{\iota}, \kappa_1), (\kappa_2, \ldots, \kappa_m), (x_1, \ldots, x_{n+1}), A', \mathsf{either}$ and $(s_1, \ldots, s_n, t_1)$ and $(t_2, \ldots, t_m)$. gives an element $(\overrightarrow{u}, \varphi) \in A'$ and γ such that $\gamma(x_i) = s_i$ for $1 \leq i \leq n$ and $\gamma(x_{n+1}) = t_1$ and each $t_{i+1} = u_i\gamma$, and $\varphi\gamma$ is a valid ground logical constraint. Now, $(\overrightarrow{u}, \varphi)$ can be written as $((w_2, \ldots, w_n), \varphi'[w_1 := x_{n+1}])$ for some $(\overrightarrow{w}, \varphi') \in A$. So let δ be the substitution $\gamma \cup [w_1 := \gamma(x_{n+1})]$. Noting that by linearity w_1 cannot occur in the other w_i, and that $\varphi'\delta = \varphi\gamma$ because x_{n+1} does not occur in φ', each further $\delta(x_i) = \gamma(x_i) = s_i$ and $t_1 = \gamma(x_{n+1}) = \delta(w_1) = w_1\delta$ and $t_i = u_i\gamma = w_i\delta$ for larger i.

If $m > 0$ and $b = \mathsf{term}$, then we can assume that t_1 is not a value. If all $(\overrightarrow{u}, \varphi) \in A$ have a variable for u_1, then we use the induction hypothesis and find a suitable element $((u_2, \ldots, u_m), \varphi) \in A'$ and substitution γ with $\gamma(x_i) = s_i$ for all $1 \leq i \leq n$, such that each $t_i = u_i\gamma$ ($i > 1$) and $\varphi\gamma$ is a valid ground logical constraint. Choose $\delta := \gamma \cup [u_1 := t_1]$ (this is safe by linearity). The same requirements are satisfied, and also $t_1 = u_1\delta$!

Finally, suppose $m > 0$ and $b = \mathsf{term}$ and A has some element $(\overrightarrow{u}, \varphi)$ where u_1 is not a variable; by assumption it is also not a value. By the conditions, we may assume that always $u_1 \notin Var(\varphi)$. Let $f_1, \ldots, f_k$ be all constructors in $\Sigma_{terms} \setminus \mathcal{V}al$ with output sort κ_1. Since also t_1 is not a value, but is a ground constructor term, it can only have the form $f_p(w_1, \ldots, w_k)$ for some $p, \overrightarrow{w}$. Observing that $\boldsymbol{OK}(\overrightarrow{x}, B_p, \mathsf{either})$ must hold, we use the induction hypothesis, for $\overrightarrow{x}$, $\overrightarrow{s}$ and $(w_1, \ldots, w_k, t_2, \ldots, t_m)$, and find both a suitable tuple $((q_1, \ldots, q_k, u_2, \ldots, u_m), \varphi) \in B_p$ and a substitution γ which respects φ, maps $\overrightarrow{x}$ to $\overrightarrow{s}$ and has $q_i\gamma = w_i$ and $u_j\gamma = t_j$ for all i, j.

By definition of B_p, we have $((u_1, \ldots, u_m), \varphi) \in A$ for some u_1 which is either a variable (in which case all q_i are fresh variables), or $u_1 = f_p(q_1, \ldots, q_k)$. In the case of a variable, u_1 cannot occur in any of the other u_i by the linearity requirement, nor in φ by the conditions. Thus, we can safely assume that u_1 does not occur in the domain of γ, and choose $\delta := \gamma \cup [u_1 := f(w_1, \ldots, w_k)]$. Then $\varphi\delta = \varphi\gamma$ is still a valid ground constraint, each $s_i = x_i\gamma = x_i\delta$ and for $i > 1$ also $t_i = u_i\gamma = u_i\delta$. Finally, $t_1 = f_p(w_1, \ldots, w_k) = u_1\delta$ as required. In the alternative case that $u_1 = f_p(q_1, \ldots, q_k)$, we observe that γ already suffices: each $s_i = x_i\gamma$, for $i > 1$ we have $t_i = u_i\gamma$ and $t_1 = f_p(w_1, \ldots, w_k) = f_p(q_1\gamma, \ldots, q_k\gamma) = u_1\gamma$. □

With this, we can easily reach our main result:

THEOREM 5.2. *A left-linear and left-value-free constructor-sound LCTRS with rules $\mathcal{R}$ is quasi-reductive if for all defined and calculation symbols f: $\boldsymbol{OK}((), A_f, \mathsf{either})$ holds, where $A_f := \{(\overrightarrow{\ell}, \varphi) \mid \ell \to r\ [\varphi] \in \mathcal{R} \cup \mathcal{R}_{\mathtt{calc}} \wedge \ell = f(\overrightarrow{\ell})\}$.*

PROOF. By Lemma 3.2 it suffices to prove that all terms of the form $f(s_1, \ldots, s_n)$ can be reduced, where f is a defined or calculation symbol and all s_i are constructor terms. This holds if there is a rule $f(\ell_1, \ldots, \ell_n) \to r\ [\varphi] \in \mathcal{R} \cup \mathcal{R}_{\mathtt{calc}}$ and a substitution γ such that each $s_i = \ell_i\gamma$ and $\varphi\gamma$ is a satisfiable constraint. By Lemma 5.1 (which we can apply because left-hand sides of rules are linear and value-free), that is exactly the case if $\boldsymbol{OK}((), A_f, \mathsf{either}) = \mathtt{true}$! □

6. CONCLUSIONS

In this paper, we have given an algorithm to prove quasi-reductivity of LCTRSs, whose core idea is to identify missing cases in the rules. Although we needed to impose certain restrictions to use this algorithm, these restrictions seem very reasonable.

Although we have not proved so here, we believe that our method is not only sound, but also complete for the class of left-linear, left-value-free constructor-sound LCTRSs. We intend to explore this in future work.

The method presented in this paper has been fully implemented in our tool Ctrl, which is available at

http://cl-informatik.uibk.ac.at/software/ctrl/

References

Franz Baader and Tobias Nipkow. 1998. *Term Rewriting and All That*. Cambridge University Press.

Stephan Falke and Deepak Kapur. 2012. Rewriting Induction + Linear Arithmetic = Decision Procedure. In *IJCAR '12*. 241–255.

Cynthia Kop and Naoki Nishida. 2013. Term Rewriting with Logical Constraints. In *FroCoS '13*. 343–358.

Cynthia Kop and Naoki Nishida. 2014. Automatic Constrained Rewriting Induction Towards Verifying Procedural Programs. In *APLAS '14*. 334–353.